

General limit value in zero-sum stochastic games

Bruno Ziliotto

the date of receipt and acceptance should be inserted later

Abstract [Bewley and Kohlberg \[1976\]](#) and [Mertens and Neyman \[1981\]](#) have respectively proved the existence of the asymptotic value and the uniform value in zero-sum stochastic games with finite state space and finite action sets. In their work, the total payoff in a stochastic game is defined either as a Cesaro mean or an Abel mean of the stage payoffs. The contribution of this paper is twofold: first, it generalizes the result of [Bewley and Kohlberg \[1976\]](#) to a more general class of payoff evaluations, and it proves with an example that this new result is tight. It also investigates the particular case of absorbing games. Second, for the uniform approach of Mertens and Neyman, this paper provides an example of absorbing game to demonstrate that there is no natural way to generalize their result to a wider class of payoff evaluations.

Keywords Stochastic games · Weighted payoffs · Asymptotic value · Shapley operator · Uniform value.

Notations

- The notation “ $X := Y$ ” means “ X is defined by the expression Y ”.
- If $A \subset B$, the complementary of A in B is denoted by $B \setminus A$.
- The set of non-negative integers is denoted by \mathbb{N} , and $\mathbb{N}^* := \mathbb{N} \setminus \{0\}$.
- The set of real numbers is denoted by \mathbb{R} , and the set of strictly positive real numbers is denoted by \mathbb{R}_+ .
- If (C, \mathcal{C}) is a measurable space, we denote by $\Delta(C)$ the set of probability measures on C . We call δ_c the Dirac measure at $c \in C$. If $C_0 \subset C$ is a finite set and $(\alpha_c)_{c \in C_0} \in \Delta(C_0)$, then $\sum_{c \in C_0} \alpha_c \delta_c$ is denoted by $\sum_{c \in C_0} \alpha_c \cdot c$.

Bruno Ziliotto
CEREMADE, Université Paris Dauphine, Place du Maréchal de Lattre de Tassigny, 75016
Paris, France.
E-mail: ziliotto@math.cnrs.fr

Introduction

Zero-sum stochastic games were introduced by [Shapley \[1953\]](#). In this model, two players repeatedly play a zero-sum game, which depends on the state of nature. At each stage, a new state of nature is drawn from a distribution based on the actions of players and the state of the previous stage. The state of nature is announced to both players, along with the actions of the previous stage. Unless mentioned explicitly, we consider *finite* stochastic games: the state space and the action sets are assumed to be finite.

There are several ways to evaluate the payoff in a stochastic game. For $n \in \mathbb{N}^*$, the payoff in the n -stage game is the Cesaro mean $\frac{1}{n} \sum_{m=1}^n g_m$, where g_m is the payoff at stage $m \geq 1$. For $\lambda \in (0, 1]$, the payoff in the λ -discounted game is the Abel mean $\sum_{m \geq 1} \lambda(1 - \lambda)^{m-1} g_m$.

Two main approaches are used to understand the properties of stochastic games with long duration :

- The asymptotic approach aims at determining if the value v_n of the n -stage game and the value v_λ of the λ -discounted game converge respectively when n goes to infinity and λ goes to 0. [Bewley and Kohlberg \[1976\]](#) have proved that in finite stochastic games, (v_n) and (v_λ) converge to the same limit. This result cannot be extended to stochastic games with compact action sets (see [Vigeral \[2013\]](#)). Neither can it be extended to the case of public imperfect observation of the state of nature (see [Ziliotto \[2013\]](#)).
- The uniform approach analyzes the existence of strategies that are approximately optimal in any n -stage game and λ -discounted game, provided that n is big enough and λ is small enough. When this is the case, the stochastic game is said to have a *uniform value*. [Mertens and Neyman \[1981\]](#) have shown that finite stochastic games have a uniform value. Note that the existence of the uniform value implies the existence of the asymptotic value.

In this paper we investigate these two approaches, when payoffs are not restricted to be Cesaro means or Abel means of stage payoffs. As in [Cardaliaguet et al. \[2012\]](#), if $\pi := (\pi_m)_{m \geq 1} \in \Delta(\mathbb{N}^*)$ is a sequence of weights, the payoff in the π -weighted game is defined as the weighted sum $\sum_{m \geq 1} \pi_m g_m$. Intuitively, a π -weighted game with long duration corresponds to the case where the $(\pi_m)_{m \geq 1}$ are close to 0 (but still summing to one), but there are many different ways to define the convergence of π to 0. Once a criterion of convergence is defined, the asymptotic approach consists in determining whether or not the value v_π of the π -weighted game converges when π goes to 0. When this is the case, the game is said to have a *general asymptotic value* (with respect to the chosen criterion). Likewise, the uniform approach deals with the existence of strategies that are approximately optimal in any π -weighted game, with π small enough. When this is the case, the game is said to have a *general uniform value* (with respect to the chosen criterion). Two main results can be found in literature:

- If $(\pi_m)_{m \geq 1} \in \Delta(\mathbb{N}^*)$ is decreasing with respect to m , and if the criterion of convergence is: π_1 goes to 0, then finite stochastic games have a general

uniform value (and thus a general asymptotic value). This result stems from the existence of the uniform value, established by [Mertens and Neyman \[1981\]](#), and from Theorem 1 and Remark (4) in [Neyman and Sorin \[2010\]](#).

- [Renault and Venel \[2012\]](#) examine payoff weights that are not necessarily decreasing with respect to time, and consider the *impatience* $I_1(\pi) := \sum_{m \geq 1} |\pi_{m+1} - \pi_m|$ of π (see [\[Sorin, 2002, section 5.7\]](#)). They investigate the limit behavior of finite stochastic games with one Player (Markov Decision Processes) and finite POMDP (Markov Decision Processes with Partial Observation), when $I_1(\pi)$ goes to 0. In this framework, they show the existence of the general uniform value. Note that if π is decreasing and π_1 goes to 0, then $I_1(\pi) = \pi_1$ goes to 0. Thus, for MDPs, this second result is more general than the first one.

In this paper, we also define a criterion on the convergence of π to 0 under which the general asymptotic value exists in stochastic games. For the asymptotic approach, our theorem generalizes the two aforementioned results. In addition, we provide an example which shows first that our result is tight, and second that the result of [Renault and Venel \[2012\]](#) cannot be extended to the Two-Player Case. We also show that for absorbing games with compact action sets and separately continuous transition and payoff functions, a sufficient condition under which (v_π) converges is that $\sup_{m \geq 1} \pi_m$ goes to 0 (when the action sets are finite, a sketch of proof for this last result is written in [Cardaliaguet et al. \[2012\]](#)). As for the uniform approach, we provide an example of absorbing game which shows that there is no natural way to relax the decreasing assumption on the weights.

The paper is organized as follows. Section 1 presents the model of stochastic games and some basic concepts. Section 2 deals with the asymptotic approach, and Section 3 presents the uniform approach.

1 Generalities

1.1 Model of stochastic game

A stochastic game Γ is defined by:

- A state space K ,
- An action set I (resp. J) for Player 1 (resp. 2),
- A payoff function $g : K \times I \times J \rightarrow [0, 1]$,
- A transition function $q : K \times I \times J \rightarrow \Delta(K)$.

Except in Subsection 2.2, we assume that K, I, J are (nonempty) finite sets. The initial state is $k_1 \in K$, and the stochastic game Γ^{k_1} which starts in k_1 proceeds as follows. At each stage $m \geq 1$, both players choose simultaneously and independently an action, $i_m \in I$ (resp. $j_m \in J$) for Player 1 (resp. 2). The payoff at stage m is $g_m := g(k_m, i_m, j_m)$. The state k_{m+1} of stage $m + 1$ is drawn from the probability distribution $q(k_m, i_m, j_m)$. Then (k_{m+1}, i_m, j_m) is

publicly announced to both players.

The set of all possible histories before stage m is $H_m := (K \times I \times J)^{m-1} \times K$. A *behavioral strategy* for Player 1 (resp. 2) is a mapping $\sigma : \cup_{m \geq 1} H_m \rightarrow \Delta(I)$ (resp. $\tau : \cup_{m \geq 1} H_m \rightarrow \Delta(J)$). The set of all behavioral strategies for Player 1 (resp. 2) is denoted by Σ (resp. \mathcal{T}).

A *pure strategy* for Player 1 (resp. 2) is a mapping $\sigma : \cup_{m \geq 1} H_m \rightarrow I$ (resp. $\tau : \cup_{m \geq 1} H_m \rightarrow J$).

A *Markov strategy* is a strategy that depends only on the current stage and state. A Markov strategy for Player 1 (resp. 2) can be assimilated to a mapping from $\mathbb{N}^* \times K$ to $\Delta(I)$ (resp. $\Delta(J)$).

A *stationary strategy* is a strategy that depends only on the current state. A stationary strategy for Player 1 (resp. 2) can be assimilated to a mapping from K to $\Delta(I)$ (resp. $\Delta(J)$).

The set of infinite plays of the game is $H_\infty := (K \times I \times J)^{\mathbb{N}^*}$, and is equipped with the σ -algebra generated by cylinders. A triple $(k_1, \sigma, \tau) \in K \times \Sigma \times \mathcal{T}$ induces a unique probability measure on H_∞ , denoted by $\mathbb{P}_{\sigma, \tau}^{k_1}$ (see [Sorin, 2002, Appendix D]). Let $\pi \in \Delta(\mathbb{N}^*)$ such that $\sum_{m \geq 1} \pi_m = 1$. The π -*weighted game* $\Gamma_\pi^{k_1}$ is the game defined by its normal form $(\Sigma, \mathcal{T}, \gamma_\pi^{k_1})$, where

$$\gamma_\pi^{k_1}(\sigma, \tau) := \mathbb{E}_{\sigma, \tau}^{k_1} \left(\sum_{m \geq 1} \pi_m g_m \right).$$

By the minmax theorem (see [Sorin, 2002, Appendix A.5]), the game $\Gamma_\pi^{k_1}$ has a value, denoted by $v_\pi(k_1)$:

$$v_\pi(k_1) = \max_{\sigma \in \Sigma} \min_{\tau \in \mathcal{T}} \gamma_\pi^{k_1}(\sigma, \tau) = \min_{\tau \in \mathcal{T}} \max_{\sigma \in \Sigma} \gamma_\pi^{k_1}(\sigma, \tau).$$

When for some $n \in \mathbb{N}^*$, $\pi_m = n^{-1} 1_{m \leq n}$ for every $m \in \mathbb{N}^*$, the game $\Gamma_n := \Gamma_\pi$ is called the n -*stage game*, and its payoff function is denoted by γ_n . When for some $\lambda \in (0, 1]$, $\pi_m = \lambda(1 - \lambda)^{m-1}$ for every $m \in \mathbb{N}^*$, the game $\Gamma_\lambda := \Gamma_\pi$ is called the λ -*discounted game*, and its payoff function is denoted by γ_λ .

1.2 Two results in the literature

Let us fix a stochastic game Γ . Two standard definitions are recalled below:

Definition 1 The stochastic game Γ has an *asymptotic value* if the sequences (v_λ) and (v_n) converge to the same limit, when respectively λ goes to 0 and n goes to infinity.

Definition 2 The stochastic game Γ has a *uniform value* $v_\infty : K \rightarrow [0, 1]$ if for all $k_1 \in K$, for all $\epsilon > 0$, there exists $(\sigma^*, \tau^*) \in \Sigma \times \mathcal{T}$ and $\bar{n} \in \mathbb{N}^*$, such that for all $n \geq \bar{n}$ and $(\sigma, \tau) \in \Sigma \times \mathcal{T}$, we have

$$\gamma_n^{k_1}(\sigma^*, \tau) \geq v_\infty(k_1) - \epsilon \quad \text{and} \quad \gamma_n^{k_1}(\sigma, \tau^*) \leq v_\infty(k_1) + \epsilon.$$

Bewley and Kohlberg [1976] have proved that Γ has an asymptotic value, and Mertens and Neyman [1981] have generalized this result in the following way:

Theorem 1 *The stochastic game Γ has a uniform value v_∞ . In particular, Γ has an asymptotic value, and (v_λ) and (v_n) converge to v_∞ , when respectively λ goes to 0 and n goes to infinity.*

This theorem shows that both players have strategies that are approximately optimal in any long game $\Gamma_n^{k_1}$. The existence of a stronger notion of uniform value is then straightforward (see Theorem 1 and Remark (4) in Neyman and Sorin [2010]): players have strategies that are approximately optimal in any game $\Gamma_\pi^{k_1}$ with $\pi = (\pi_m)_{m \geq 1}$ decreasing with respect to m and π_1 sufficiently small. For completeness, we give a sketch of the proof of this corollary.

Corollary 1 *For all $k_1 \in K$, for all $\epsilon > 0$, there exists $(\sigma^*, \tau^*) \in \Sigma \times \mathcal{T}$ and $\alpha > 0$ such that for all $\pi = (\pi_m)_{m \geq 1} \in \Delta(\mathbb{N}^*)$ decreasing with respect to m , and that satisfies $I_\infty(\pi) := \sup_{m \geq 1} \pi_m = \pi_1 \leq \alpha$, we have for all $(\sigma, \tau) \in \Sigma \times \mathcal{T}$*

$$\gamma_\pi^{k_1}(\sigma^*, \tau) \geq v_\infty(k_1) - \epsilon \quad \text{and} \quad \gamma_\pi^{k_1}(\sigma, \tau^*) \leq v_\infty(k_1) + \epsilon.$$

In particular, (v_π) converges to v_∞ when π is decreasing and $I_\infty(\pi)$ goes to 0.

Proof (Sketch) If $\pi \in \Delta(\mathbb{N}^*)$ is decreasing, then the π -weighted payoff is a convex combination of Cesaro-mean payoffs:

$$\sum_{m \geq 1} \pi_m g_m = \sum_{m \geq 1} m(\pi_m - \pi_{m+1}) \frac{1}{m} \sum_{l=1}^m g_l.$$

The proof of the corollary follows from this equality. \square

In the One-Player Case, by a particular case of Theorem 3.19 in Renault and Venel [2012], this result can be extended to a wider class of weights, in the following way:

Theorem 2 *Assume that Γ is a Markov decision process, that is, the functions q and g do not depend on the action of Player 2. Then Γ has a general uniform value: for all $k_1 \in K$, for all $\epsilon > 0$, there exists $\sigma^* \in \Sigma$ and $\alpha > 0$, such that for all $\pi \in \Delta(\mathbb{N}^*)$ that satisfies*

$$I_1(\pi) := \sum_{m \geq 1} |\pi_{m+1} - \pi_m| \leq \alpha, \text{ we have}$$

$$v_\infty(k_1) - \epsilon \leq \gamma_\pi^{k_1}(\sigma^*) \leq v_\infty(k_1) + \epsilon.$$

In particular, (v_π) converges to v_∞ when $I_1(\pi)$ goes to 0.

In the next section, we study the asymptotic approach and investigate whether we can also relax the decreasing assumption in Corollary 1 in the Two-Player Case.

2 Asymptotic approach

2.1 A criterion for the convergence of (v_π)

The first obvious point is that if one removes the decreasing assumption in Corollary 1 and only assumes that $I_\infty(\pi) := \sup_{m \geq 1} \pi_m$ goes to zero, (v_π) does not necessarily converge. Indeed consider the Markov chain which oscillates deterministically between two states, one with payoff 1, the other one with payoff 0. Consider two sequences of weights, one which puts weight on even stages and one which puts weight on odd stages. The difference between these two payoff evaluations is always equal to 1. Thus the condition $I_\infty(\pi) \rightarrow 0$ is not a sufficient condition to obtain the convergence of (v_π) . Let us now provide a more restrictive criterion under which (v_π) converges.

Definition 3 Let $\pi \in \Delta(\mathbb{N}^*)$ and $p \in (0, +\infty]$. The p -impatience of π is the quantity $I_p(\pi) \in (0, +\infty]$ defined by

$$I_p(\pi) := \begin{cases} \sum_{m \geq 1} |(\pi_{m+1})^p - (\pi_m)^p| & \text{if } p < \infty, \\ \sup_{m \geq 1} \pi_m & \text{if } p = \infty. \end{cases}$$

When $I_p(\pi)$ is small, it means that players are very patient. When in addition $p < \infty$, it means that the variations of π with respect to m are small.

Proposition 1 Let $\pi \in \Delta(\mathbb{N}^*)$ and $p, p' \in \mathbb{R}_+^*$, such that $p \leq p'$. Then

- $I_{p'}(\pi) \leq (p'/p)I_p(\pi)$,
- $I_\infty(\pi) \leq (I_p(\pi))^{1/p}$.

Proof Let $m \in \mathbb{N}^*$ and $q := p'/p$. The Mean Value Theorem implies that

$$\left| (\pi_{m+1})^{p'} - (\pi_m)^{p'} \right| = \left| [(\pi_{m+1})^p]^q - [(\pi_m)^p]^q \right| \leq q |(\pi_{m+1})^p - (\pi_m)^p|,$$

and it yields: $I_{p'}(\pi) \leq qI_p(\pi)$. As for the second inequality, we have

$$(\pi_m)^p = \sum_{m' \geq m} [(\pi_{m'})^p - (\pi_{m'+1})^p] \leq I_p(\pi),$$

and it yields: $I_\infty(\pi) \leq (I_p(\pi))^{1/p}$. \square

Remark 1 When $(\pi_m)_{m \geq 1}$ is decreasing, for all $p \in \mathbb{R}_+^*$, we have $I_p(\pi) = (\pi_1)^p$. Consequently, given $p, p' \in \mathbb{R}_+^*$ such that $p \leq p'$, there does not exist a real number $C(p, p') > 0$ such that for all $\pi \in \Delta(\mathbb{N}^*)$, $I_{p'}(\pi) \geq C(p, p')I_p(\pi)$.

Let us fix a stochastic game Γ , and let $v_\infty = \lim_{\lambda \rightarrow 0} v_\lambda = \lim_{n \rightarrow +\infty} v_n$ be its uniform value. For f a real-valued function, denote by $\|f\|_\infty$ the supremum of f .

Definition 4 Let $p \in (0, +\infty]$. The stochastic game Γ has a *p-asymptotic value* if for all $\epsilon > 0$, there exists $\alpha > 0$ such that for all $\pi \in \Delta(\mathbb{N}^*)$ verifying $I_p(\pi) \leq \alpha$, we have $\|v_\pi - v_\infty\|_\infty \leq \epsilon$.

Remark 2

- If for some $p' \in (0, +\infty]$, the game Γ has a p' -asymptotic value, it has a p -asymptotic value for all $p \leq p'$. It results directly from Proposition 1.
- By Theorem 2, any Markov decision process has a 1-asymptotic value.
- Finite absorbing games have an ∞ -asymptotic value (see Cardaliaguet et al. [2012]).
- The Markov chain described at the beginning of this subsection has no p -asymptotic value for all $p > 1$.

Recall that (v_λ) can be expanded in Puiseux series (see Bewley and Kohlberg [1976]): there exists $\beta > 0$, $M \in \mathbb{N}^*$ and $r_m \in \mathbb{R}^K$ such that for all $k \in K$ and $\lambda \in [0, \beta)$

$$v_\lambda(k) = \sum_{m \geq 0} r_m(k) \lambda^{\frac{m}{M}}, \quad (1)$$

with the convention $v_0 := v_\infty$.

Definition 5 Let $m_0 = \inf \{m \geq 1 \mid r_m \neq 0\}$. The quantity $s := m_0/M \in [0, +\infty]$ is called the *order* of Γ .

Note that if $s < +\infty$, there exists $C > 0$ such that for all $(\lambda, \lambda') \in [0, \beta)^2$, we have

$$\|v_\lambda - v_{\lambda'}\|_\infty \leq C |\lambda^s - \lambda'^s|. \quad (2)$$

If A is a finite set, the cardinal of A is denoted by $\text{Card } A$. By Remark 3 in Oliu-Barton [2014], we have

$$s \geq (\text{Card } K \text{ Card } I)^{-\sqrt{\text{Card } K \text{ Card } I}}.$$

Now we can state our main theorem.

Theorem 3 *The stochastic game Γ has a s-asymptotic value. In particular, if $p \in \mathbb{R}_+^*$ is smaller or equal to $(\text{Card } K \text{ Card } I)^{-\sqrt{\text{Card } K \text{ Card } I}}$, then Γ has a p-asymptotic value.*

Proof Neyman has shown that in a stochastic game the convergence of (v_n) can be deduced from the Shapley equation and the fact that (v_λ) is absolutely continuous with respect to λ (see [Sorin, 2002, Theorem C.8, p. 177]). We use similar tools.

Let $\pi \in \Delta(\mathbb{N}^*)$ and $r \in \mathbb{N}$ such that there exists $m \geq r + 2$, $\pi_m \neq 0$. A sequence of weights $\pi^r \in \Delta(\mathbb{N}^*)$ is defined in the following way: for $m \in \mathbb{N}^*$,

$$\pi_m^r := \begin{cases} \frac{\pi_{m+r}}{\sum_{m' \geq r+1} \pi_{m'}} & \text{if } \pi_{m+r} \neq 0, \\ 0 & \text{if } \pi_{m+r} = 0. \end{cases}$$

Let $\lambda_r := \pi_1^r$. Let $k \in K$. Shapley equations yield (see [Cardaliaguet et al. \[2012\]](#)):

$$v_{\pi^r}(k) = \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} \{ \lambda_r g(k, x, y) + (1 - \lambda_r) \mathbb{E}_{x,y}^k(v_{\pi^{r+1}}) \} \quad (3)$$

$$= \min_{y \in \Delta(J)} \max_{x \in \Delta(I)} \{ \lambda_r g(k, x, y) + (1 - \lambda_r) \mathbb{E}_{x,y}^k(v_{\pi^{r+1}}) \} \quad (4)$$

and

$$v_{\lambda_r}(k) = \max_{x \in \Delta(I)} \min_{y \in \Delta(J)} \{ \lambda_r g(k, x, y) + (1 - \lambda_r) \mathbb{E}_{x,y}^k(v_{\lambda_r}) \} \quad (5)$$

$$= \min_{y \in \Delta(J)} \max_{x \in \Delta(I)} \{ \lambda_r g(k, x, y) + (1 - \lambda_r) \mathbb{E}_{x,y}^k(v_{\lambda_r}) \}, \quad (6)$$

where

$$\mathbb{E}_{x,y}^k(f) := \sum_{(k', i, j) \in K \times I \times J} x(i) y(j) q(k, i, j)(k') f(k') \quad (7)$$

and

$$g(k, x, y) := \sum_{(i, j) \in I \times J} x(i) y(j) g(k, i, j). \quad (8)$$

Note that these equations also hold when $\lambda_r = 0$, with the convention $v_0 := v_\infty$. Let $x \in \Delta(I)$ be optimal in (3) and $y \in \Delta(J)$ be optimal in (6). We have

$$v_{\pi^r}(k) \leq \lambda_r g(k, x, y) + (1 - \lambda_r) \mathbb{E}_{x,y}^k(v_{\pi^{r+1}}) \quad (9)$$

and

$$v_{\lambda_r}(k) \geq \lambda_r g(k, x, y) + (1 - \lambda_r) \mathbb{E}_{x,y}^k(v_{\lambda_r}). \quad (10)$$

Combining the two inequalities yields:

$$v_{\pi^r}(k) - v_{\lambda_r}(k) \leq (1 - \lambda_r) \|v_{\pi^{r+1}} - v_{\lambda_r}\|_\infty.$$

Symmetrically, if $x' \in \Delta(I)$ is optimal in (4) and $y' \in \Delta(J)$ is optimal in (5), then

$$v_{\lambda_r}(k) - v_{\pi^r}(k) \leq (1 - \lambda_r) \|v_{\pi^{r+1}} - v_{\lambda_r}\|_\infty,$$

and thus

$$\|v_{\pi^r} - v_{\lambda_r}\|_\infty \leq (1 - \lambda_r) \|v_{\pi^{r+1}} - v_{\lambda_r}\|_\infty. \quad (11)$$

Let $\Pi_r := \prod_{r'=0}^{r-1} (1 - \lambda_{r'}) = \sum_{m \geq r+1} \pi_m$. Note that $\lim_{r \rightarrow +\infty} \Pi_r = 0$. The last inequality yields:

$$\Pi_r \|v_{\pi^r} - v_{\lambda_r}\|_\infty \leq \Pi_{r+1} \|v_{\pi^{r+1}} - v_{\lambda_{r+1}}\|_\infty + \Pi_{r+1} \|v_{\lambda_{r+1}} - v_{\lambda_r}\|_\infty.$$

Let $N \in \mathbb{N}^*$ such that there exists $m \geq N+1$, $\pi_m \neq 0$. Summing this inequality over $r \in \{0, 1, \dots, N-1\}$ yields:

$$\|v_\pi - v_{\lambda_0}\|_\infty \leq \Pi_N \|v_{\pi^N} - v_{\lambda_N}\|_\infty + \sum_{r=1}^N \Pi_r \|v_{\lambda_r} - v_{\lambda_{r-1}}\|_\infty. \quad (12)$$

Let $\epsilon \in (0, 1)$. Let $N_0 := \max \{N \geq 1 \mid \Pi_N \geq \epsilon\}$. We have $\Pi_{N_0} \leq \epsilon + I_\infty(\pi)$. For $N = N_0$, inequality (12) writes:

$$\|v_\pi - v_{\lambda_0}\|_\infty \leq \epsilon + I_\infty(\pi) + \sum_{r=1}^{N_0} \|v_{\lambda_r} - v_{\lambda_{r-1}}\|_\infty. \quad (13)$$

Assume that $I_\infty(\pi) < \epsilon\beta$ (see equation (1) for the definition of β). Thus $\lambda_r \leq I_\infty(\pi)/\Pi_{N_0} < \beta$ for all $r \in \{0, 1, \dots, N_0\}$. If $s = \infty$, then $\sum_{r=1}^{N_0} \|v_{\lambda_r} - v_{\lambda_{r-1}}\|_\infty = 0$, and the last inequality proves that Γ has an ∞ -asymptotic value. Assume now that $s < \infty$. Using (2), let us majorize the term on the right in inequality (13):

$$\sum_{r=1}^{N_0} \|v_{\lambda_r} - v_{\lambda_{r-1}}\|_\infty \leq C \sum_{r=1}^{N_0} |(\lambda_r)^s - (\lambda_{r-1})^s|.$$

Let $r \in \{1, 2, \dots, N_0\}$. The quantity $|(\lambda_r)^s - (\lambda_{r-1})^s|$ is smaller than

$$\left| \frac{(\pi_{r+1})^s}{\left(\sum_{m \geq r+1} \pi_m\right)^s} - \frac{(\pi_r)^s}{\left(\sum_{m \geq r+1} \pi_m\right)^s} \right| + \left| \frac{(\pi_r)^s}{\left(\sum_{m \geq r+1} \pi_m\right)^s} - \frac{(\pi_r)^s}{\left(\sum_{m \geq r} \pi_m\right)^s} \right|.$$

By definition of N_0 , we have

$$\sum_{m \geq r} \pi_m \geq \sum_{m \geq r+1} \pi_m \geq \epsilon.$$

Therefore we can majorize the term on the left by $\epsilon^{-s} |(\pi_{r+1})^s - (\pi_r)^s|$. As for the term on the right, by the Mean Value theorem we have

$$\begin{aligned} \left(\sum_{m \geq r+1} \pi_m \right)^{-s} - \left(\sum_{m \geq r} \pi_m \right)^{-s} &\leq s \left(\sum_{m \geq r+1} \pi_m \right)^{-1-s} \pi_r \\ &\leq s \epsilon^{-1-s} \pi_r. \end{aligned}$$

Finally we have

$$\begin{aligned} \sum_{r \geq 1} |(\lambda_r)^s - (\lambda_{r-1})^s| &\leq \sum_{r \geq 1} (\epsilon^{-s} |(\pi_{r+1})^s - (\pi_r)^s| + s \epsilon^{-1-s} (\pi_r)^{1+s}) \\ &\leq \epsilon^{-s} I_s(\pi) + s \epsilon^{-1-s} I_\infty(\pi)^s \\ &\leq (\epsilon^{-s} + s \epsilon^{-1-s}) I_s(\pi). \end{aligned}$$

Plugging this into (13) gives

$$\|v_\pi - v_{\lambda_0}\|_\infty \leq \epsilon + I_s(\pi)^{1/s} + C (\epsilon^{-s} + s \epsilon^{-1-s}) I_s(\pi).$$

Thus for $I_s(\pi)$ sufficiently small, we have both $\|v_\pi - v_{\lambda_0}\|_\infty \leq \epsilon$ and $\|v_{\lambda_0} - v_\infty\|_\infty \leq \epsilon$, which concludes the proof. \square

Corollary 2 *Let $(\pi^n) \in (\Delta(\mathbb{N}^*))^{\mathbb{N}}$ such that for all $p > 0$, $\lim_{n \rightarrow +\infty} I_p(\pi^n) = 0$. Then in any stochastic game, $(v_{\pi^n})_{n \geq 0}$ converges to v_∞ .*

Proof Let Γ be a stochastic game of order $s \in (0, +\infty]$. By Theorem 3, Γ has a s -asymptotic value. By assumption, we have $\lim_{n \rightarrow +\infty} I_s(\pi^n) = 0$, thus $(v_{\pi^n})_{n \geq 0}$ converges to v_∞ . \square

The following remarks show that for the asymptotic approach, Corollary 2 is more general than Corollary 1 and Theorem 2.

Remark 3

- When $(\pi_m)_{m \geq 1}$ is decreasing, for all $p > 0$, $I_p(\pi) = (\pi_1)^p$. According to Corollary 2, (v_π) converges when π_1 goes to 0 (compare with the asymptotic approach in Corollary 1).
- When $s = 1$, the mapping $\lambda \rightarrow v_\lambda$ is Lipschitz. For instance, this is the case when Γ is a Markov decision process: see [Sorin, 2002, Chapter 5, Proposition 5.20]. By Theorem 3, Γ has a 1-asymptotic value (compare with the asymptotic approach in Theorem 2).
- For $(l, n) \in \mathbb{N} \times \mathbb{N}^*$, let $\pi^{l,n} = n^{-1} 1_{l+1 \leq m \leq l+n}$. The $(\pi_{l,n})$ are non-monotonic sequences, thus Corollary 1 does not apply. Nevertheless, $I_s(\pi^{l,n}) = 2n^{-s}$. Consequently, for any $\epsilon > 0$, there exists $\bar{n} \in \mathbb{N}^*$ such that for all $n \geq \bar{n}$, for all $l \in \mathbb{N}$, $I_s(\pi^{l,n}) \leq \epsilon$. By Theorem 3, Γ has a s -asymptotic value, and we deduce that

$$\lim_{n \rightarrow +\infty} \sup_{l \in \mathbb{N}} \|v_{\pi^{l,n}} - v_\infty\|_\infty = 0.$$

2.2 Absorbing games

In this subsection, we relax the finiteness assumption on the action sets. An absorbing state is a state such that once it is reached, the game remains in this state forever, and the payoff does not depend on the actions (*absorbing payoff*). An absorbing game is a stochastic game that has at most one nonabsorbing state.

Mertens et al. [2009] have proved the existence of the uniform value in absorbing games with compact action sets and separately continuous transition and payoff functions. In particular, (v_λ) converges. Adapting the proof of the previous subsection, we prove the following proposition:

Proposition 2 *Let Γ be an absorbing game with compact action sets and separately continuous transition and payoff functions. Then, Γ has an ∞ -asymptotic value.*

Remark 4 For finite I and J , this result was stated in Cardaliaguet et al. [2012], with a sketch of proof. Here, we provide a complete and simpler demonstration, which holds in a more general framework.

Proof Again, we adopt the convention $v_0 := v_\infty$. The Shapley equations (3), (4), (5) and (6) still hold true for compact action sets (see Maitra and Parthasarathy [1970]). The only difference is that in (7) and (8), the sum has to be replaced by an integral.

If k^* is an absorbing state, we have $v_\pi(k^*) = v_\lambda(k^*)$, for any $\pi \in \Delta(\mathbb{N}^*)$ and $\lambda \in [0, 1]$. Let k be the only non-absorbing state of the game, and $r \in \mathbb{N}$. In the previous proof, inequalities (9) and (10) yield:

$$(v_{\pi^r} - v_{\lambda_r})(k) \leq (1 - \lambda_r)\mu_r(v_{\pi^{r+1}} - v_{\lambda_r})(k),$$

where μ_r is the probability that the game is not absorbed, when Player 1 (resp. 2) plays an optimal strategy x (resp. y) in (3) (resp. (6)). In what follows, for simplicity we omit the variable k .

Let $\Pi_r = \prod_{m=0}^{r-1} (1 - \lambda_r)\mu_r$. Relying on the same steps as in the previous proof, we get the analogous of (13), where N_0 is defined similarly:

$$(v_\pi - v_{\lambda_0}) \leq (\epsilon + I_\infty(\pi)) + \sum_{r=1}^{N_0} \Pi_r (v_{\lambda_r} - v_{\lambda_{r-1}}). \quad (14)$$

Let us majorize the right-hand side. The sequence $(\Pi_r)_{r \geq 0}$ is decreasing. Moreover, for all $r \in \{1, 2, \dots, N_0\}$, we have

$$\sum_{m \geq r+1} \pi_m \geq \Pi_r \geq \epsilon.$$

Hence, $\lambda_r \in [0, I_\infty(\pi)/\epsilon]$. Let $V := \sup_{\lambda \in [0, I_\infty(\pi)/\epsilon]} v_\lambda$. We have:

$$\begin{aligned} \sum_{r=1}^{N_0} \Pi_r (v_{\lambda_r} - v_{\lambda_{r-1}}) &= \sum_{r=1}^{N_0} \Pi_r v_{\lambda_r} - \sum_{r=0}^{N_0-1} \Pi_{r+1} v_{\lambda_r} \\ &= \sum_{r=1}^{N_0-1} (\Pi_r - \Pi_{r+1}) v_{\lambda_r} + \Pi_{N_0} v_{\lambda_{N_0}} - \Pi_1 v_{\lambda_0} \\ &\leq V \sum_{r=1}^{N_0-1} (\Pi_r - \Pi_{r+1}) + \Pi_{N_0} v_{\lambda_{N_0}} - \Pi_1 v_{\lambda_0} \\ &\leq V(\Pi_1 - \Pi_{N_0}) + \Pi_{N_0} v_{\lambda_{N_0}} - \Pi_1 v_{\lambda_0} \\ &= \Pi_1 (V - v_{\lambda_0}) - \Pi_{N_0} (V - v_{\lambda_{N_0}}) \\ &\leq |V - v_{\lambda_0}| + |V - v_{\lambda_{N_0}}|. \end{aligned}$$

Because (v_λ) converges to v_∞ when λ goes to 0, the term $|V - v_\lambda|$ vanishes when λ goes to 0. Consequently, the right-hand side of the last inequality goes to 0 when $I_\infty(\pi)$ goes to 0. Together with (14), this shows that the positive part of $(v_\pi - v_{\lambda_0})$ goes to 0 when $I_\infty(\pi)$ goes to 0.

Symmetrically, one can show that the negative part of $(v_\pi - v_{\lambda_0})$ goes to 0 when $I_\infty(\pi)$ goes to 0. Hence, the proposition is proved. \square

2.3 An example

We construct a stochastic game of order $1/2$, which has no p -asymptotic value for any $p > 1/2$. First, this shows that our main result (see Theorem 3) cannot be improved, second this implies that Theorem 2 does not extend to the Two-Player Case.

Let us consider the following stochastic game Γ :

Table 1 Transition and payoff functions in state ω_1 and ω_2

ω_1	L	R	ω_2	L	R
T	1	$\overrightarrow{1}$	T	0	$\overleftarrow{0}$
M	0	0	M	0	0
B	$\overrightarrow{1}$	1^*	B	$\overleftarrow{0}$	0^*

The set of states of the game is $K = \{\omega_1, \omega_2, 1^*, 0^*\}$. The action set is $I = \{T, M, B\}$ for Player 1 and $J = \{L, R\}$ for Player 2. States 1^* and 0^* are absorbing states with absorbing payoff respectively 1 and 0. The payoff and transition functions in state ω_1 (resp. ω_2) are described by the left table (resp. the right one). The symbol $\overrightarrow{1}$ (resp. $\overleftarrow{0}$) means that the payoff is 1 (resp. 0) and the game moves on to state ω_2 (resp. ω_1). When there is no arrow or star, this means that the game remains in the same state.

In Vigeral [2013], a similar stochastic game Γ' is mentioned. The only difference is that in Γ' , Player 1 has only the two actions T and B . The uniform value v'_∞ of Γ' satisfies $v'_\infty(\omega_1) = v'_\infty(\omega_2) = 1/2$. In addition, the order of Γ' is $1/2$. Moreover, for all $\epsilon > 0$, the stationary strategy x (resp. y) for Player 1 (resp. 2) defined by $x(\omega_1) = x(\omega_2) = (1 - \sqrt{\lambda}) \cdot T + \sqrt{\lambda} \cdot B$ (resp. $y(\omega_1) = y(\omega_2) = (1 - \sqrt{\lambda}) \cdot L + \sqrt{\lambda} \cdot R$) is ϵ -optimal in Γ'_λ , for λ small enough (they are *asymptotically optimal* strategies).

In our example, in Γ_λ , the action M is dominated by T in every state. Thus for all $\lambda \in (0, 1]$, $v_\lambda = v'_\lambda$. In particular, Γ has order $1/2$, and its uniform value v_∞ satisfies $v_\infty(\omega_1) = v_\infty(\omega_2) = 1/2$. In addition, the strategy x (resp. y) is an asymptotically optimal stationary strategy for Player 1 (resp. 2) in Γ_λ .

Remark 5

- Assume that for some $\alpha > 0$, Player 1 plays $(1 - \alpha) \cdot T + \alpha \cdot B$ in state ω_2 until the state changes. Whatever Player 2 plays, Player 1 spends at most a number of stages of order α^{-1} in state ω_2 before moving to state ω_1 or 0^* , and the probability that the state goes to 0^* and not to ω_1 is at most of order α . Hence, for Player 1 there is a trade-off between staying not too long in state ω_2 , and having a low probability of being absorbed in 0^* . In view of what precedes, the optimal trade-off in Γ_λ is $\alpha \approx \sqrt{\lambda}$.

- Let $\theta \in \Delta(\mathbb{N}^*)$. At some stage m in Γ_θ , the action M may not be dominated by T in state ω_1 . Indeed, if $\theta_m = 0$, the stage payoff is 0 whatever the actions played. Thus, it is optimal for Player 1 to play M , because it makes the state remain in ω_1 . The example builds on this fact. By contrast, for any $\theta \in \Delta(\mathbb{N}^*)$, in Γ_θ , the action M is dominated by T in state ω_2 . In what follows, we build a family of strategies for Player 1 that all use action M in state ω_2 , but this is only to make the proof easier.

Theorem 4 *For all $p > 1/2$, the game Γ has no p -asymptotic value.*

The remainder of the subsection is dedicated to the proof of Theorem 4. Let us introduce the following piece of notation: given three sequences of strictly positive real numbers $(u_n)_{n \geq 1}$, $(v_n)_{n \geq 1}$ and $(w_n)_{n \geq 1}$, we write $u_n = v_n + o(w_n)$ if the sequence $([u_n - v_n]/w_n)_{n \geq 1}$ converges to 0.

To simplify the presentation, we first show that Γ has no 1-asymptotic value. Let $n \in \mathbb{N}^*$. For $l \in \{0, 1, \dots, n^3 - 1\}$, define $a_n(l) := l(n + n^5) + 1$ and $b_n(l) := l(n + n^5) + n$. Let $E_1 := \bigcup_{0 \leq l \leq n^3 - 1} \{a_n(l), a_n(l) + 1, \dots, b_n(l)\}$.

We consider the sequence $\pi^n \in \Delta(\mathbb{N}^*)$ defined by $\pi_m^n := n^{-4}$ if $m \in E_1$, and $\pi_m^n := 0$ otherwise. We have

$$I_1(\pi^n) = (2n^3 - 1)n^{-4},$$

thus $\lim_{n \rightarrow +\infty} I_1(\pi^n) = 0$. We show below that $\lim_{n \rightarrow +\infty} v_{\pi^n}(\omega_1) = 1$.

We consider the Markovian strategy $\sigma^n \in \Sigma$ for Player 1, described by the following table:

Table 2 Strategy σ^n

stage \ state	$m \in E_1$	$m \notin E_1$
$k_m = \omega_1$	$(1 - n^{-2}) \cdot T + n^{-2} \cdot B$	M
$k_m = \omega_2$	M	$(1 - n^{-4}) \cdot T + n^{-4} \cdot B$

We show that for any $\epsilon > 0$, for any n sufficiently large, σ^n guarantees the payoff $1 - \epsilon$ in $\Gamma_{\pi^n}^{\omega_1}$ for Player 1.

Let τ^n be a pure Markovian best-response to σ^n in $\Gamma_{\pi^n}^{\omega_1}$. Let Ω_n be the event

$$\Omega_n := \bigcap_{l \in \{0, 1, \dots, n^3 - 1\}} \{k_{a_n(l)} \in \{\omega_1, 1^*\}\}.$$

When the state of the game is ω_2 and $m \notin E_1$, Player 1 plays B with probability n^{-4} . By Remark 5, Player 1 spends at most a number of stages of order n^4 in ω_2 , and the state goes to 0^* with a probability at most of order n^{-4} . As a result, if for some $l \in \{0, \dots, n^3 - 1\}$ the state is in ω_2 at stage $b_n(l) + 1$, the probability that it will move to ω_1 before stage $a_n(l + 1)$ is

at least of order $1 - n^{-4}$. Once the state has moved to ω_1 , Player 1 plays M and the state remains in ω_1 until stage $a_n(l+1)$. Hence the probability that $k_{a_n(l)}$ lies in $\{\omega_1, 1^*\}$ for any l in $\{0, 1, \dots, n^3 - 1\}$ is at least of order $(1 - n^{-4})^{n^3} = 1 + o(1)$. This informal discussion provides intuition for the following lemma.

Lemma 1

$$\lim_{n \rightarrow +\infty} \mathbb{P}_{\sigma^n, \tau^n}^{\omega_1}(\Omega_n) = 1.$$

For notational convenience, in the proof of this lemma and the proof of the next proposition, for $n \in \mathbb{N}^*$, $\mathbb{P}_{\sigma^n, \tau^n}^{\omega_1}$ is denoted by \mathbb{P} and $\mathbb{E}_{\sigma^n, \tau^n}^{\omega_1}$ is denoted by \mathbb{E} .

Proof Let $n \in \mathbb{N}^*$ and $l \in \{0, \dots, n^3 - 1\}$. Let us minorize the probability $\mathbb{P}(k_{a_n(l+1)} \in \{\omega_1, 1^*\} | k_{a_n(l)} \in \{\omega_1, 1^*\})$.

First, notice that $\mathbb{P}(k_{b_n(l)+1} \neq 0^* | k_{a_n(l)} \in \{\omega_1, \omega_1^*\}) = 1$ (see Table 2). Let us now analyze how the state may evolve during the block

$\{b_n(l) + 1, b_n(l) + 2, \dots, a_n(l+1) - 1\}$, discriminating between the case $k_{b_n(l)} = 1^*$, $k_{b_n(l)} = \omega_1$, and $k_{b_n(l)} = \omega_2$:

- $\mathbb{P}(k_{a_n(l+1)} = 1^* | k_{b_n(l)+1} = 1^*) = 1$.
- If $k_{b_n(l)+1} = \omega_1$, then Player 1 will play M at each stage $m \in \{b_n(l) + 1, b_n(l) + 2, \dots, a_n(l+1) - 1\}$. Therefore, the state will remain in ω_1 :
 $\mathbb{P}(k_{a_n(l+1)} = \omega_1 | k_{b_n(l)+1} = \omega_1) = 1$.
- If $k_{b_n(l)+1} = \omega_2$, then Player 1 will play $(1 - n^{-4}) \cdot T + n^{-4} \cdot B$ as long as the state is ω_2 and $m \leq a_n(l+1) - 1$. We discriminate between two cases:
 - if Player 2 plays L as long as the state is ω_2 and $m \leq a_n(l+1) - 1$, the game will never be absorbed in 0^* , and the probability that the state will move to ω_1 before stage $a_n(l+1)$ is equal to $1 - (1 - n^{-4})^{n^5}$. If the state moves to ω_1 at some stage $m \leq a_n(l+1) - 1$, then Player 1 will play M until stage $a_n(l+1)$, thus the state will remain in ω_1 . Consequently, in this case we have

$$\mathbb{P}(k_{a_n(l+1)} = \omega_1 | k_{b_n(l)+1} = \omega_2) = 1 - (1 - n^{-4})^{n^5}.$$

- if Player 2 plays R at one stage in $\{b_n(l) + 1, b_n(l) + 2, \dots, a_n(l+1) - 1\}$, and if at the first stage he does so the state is ω_2 , then with probability $1 - n^{-4}$ the state will move to ω_1 . It will remain in ω_1 until stage $a_n(l+1)$. If the state has already switched to ω_1 before Player 2 plays R , then it will remain in ω_1 until stage $a_n(l+1)$. Therefore, in this case we have

$$\mathbb{P}(k_{a_n(l+1)} = \omega_1 | k_{b_n(l)+1} = \omega_2) \geq 1 - n^{-4}.$$

The last two subcases show that

$$\mathbb{P}(k_{a_n(l+1)} = \omega_1 | k_{b_n(l)+1} = \omega_2) \geq \min \left\{ 1 - (1 - n^{-4})^{n^5}, 1 - n^{-4} \right\}.$$

This exhaustive study shows that

$$\mathbb{P}(k_{a_n(l+1)} \in \{\omega_1, 1^*\} | k_{a_n(l)} \in \{\omega_1, 1^*\}) \geq \min \left\{ 1 - (1 - n^{-4})^{n^5}, 1 - n^{-4} \right\}.$$

We have $(1 - n^{-4})^{n^5} = o(n^{-4})$, thus for n large enough, the minimum in the above equation is reached at $1 - n^{-4}$. By induction, it yields

$$\mathbb{P}(\Omega_n) \geq \prod_{l=0}^{n^3-1} (1 - n^{-4}) = 1 + o(1),$$

and the lemma is proved. \square

Now we can prove the following proposition:

Proposition 3 *The game Γ has no 1-asymptotic value.*

Proof Let $n \in \mathbb{N}^*$. We minorize $\gamma_{\pi^n}^{\omega_1}(\sigma^n, \tau^n)$ by a quantity that goes to 1 as n goes to infinity.

The last lemma shows that with high probability, at the beginning of each block $\{a_n(l), a_n(l) + 1, \dots, b_n(l)\}$, the state is either ω_1 or 1^* . Recall that these blocks exactly correspond to the stages where the payoff weight is nonzero. Hence, to get a good payoff between stage $a_n(l)$ and stage $b_n(l)$, Player 2 should make the state move from ω_1 to ω_2 at least before stage $b_n(l)$. If Player 2 plays L at each stage $m \in \{a_n(l), a_n(l) + 1, \dots, b_n(l)\}$, with probability greater than $1 - (1 - n^{-2})^n$, the state will remain in ω_1 until stage $b_n(l)$. This probability goes to 1 as n goes to infinity, which is a bad outcome for Player 2. Thus, Player 2 should play R at some stage $m \in \{a_n(l), a_n(l) + 1, \dots, b_n(l)\}$. We show that:

- either the number of $l \in \{0, 1, \dots, n^3 - 1\}$ such that Player 2 plays at least one time R in $\{a_n(l), a_n(l) + 1, \dots, b_n(l)\}$ is small, and thus the total payoff in $I_{\pi^n}^{\omega_1}$ is close to 1,
- either the number of $l \in \{0, 1, \dots, n^3 - 1\}$ such that Player 2 plays at least one time R in $\{a_n(l), a_n(l) + 1, \dots, b_n(l)\}$ is high. In this case, with probability close to 1, the state is absorbed in 1^* very rapidly, thus the total payoff in $I_{\pi^n}^{\omega_1}$ is close to 1.

Let $n \in \mathbb{N}^*$ and $l \in \{0, 1, \dots, n^3 - 1\}$, and $\Omega_n(l)$ be the event defined by

$$\Omega_n(l) := \bigcap_{0 \leq l' \leq l} \{k_{a_n(l')} \in \{\omega_1, 1^*\}\}.$$

Note that $\Omega_n(n^3 - 1) = \Omega_n$. Let

$$M_n(l) := \{l' \in \{0, 1, \dots, l\} \mid \exists m \in \{a_n(l'), a_n(l') + 1, \dots, b_n(l')\}, \tau^n(m, \omega_1) = R\},$$

and let $\overline{M_n(l)} := \{0, 1, \dots, l\} \setminus M_n(l)$. If $l \in M_n(n^3 - 1)$, let

$$m_n(l) := \min \{m \in \{a_n(l), a_n(l) + 1, \dots, b_n(l)\} \mid \tau^n(m, \omega_1) = R\}.$$

Fix $\delta \in (0, 1]$. Let $l_n := \max \{l \in \{0, 1, \dots, n^3 - 1\} \mid \text{Card } M_n(l) \leq \delta n^3\}$. We show that between stages 1 and $b_n(l_n)$, Player 2 did not play R a sufficient number of times to impact the total payoff, and at stage $b_n(l_n) + 1$, either $l_n = n^3 - 1$ and the game is finished, or he has played too many times R , in such a way that the state has been absorbed in 1^* with high probability.

By definition of l_n , we have $\text{Card } M_n(l_n) \leq \delta n^3$, and if $l_n < n^3 - 1$, then $\text{Card } M_n(l_n) \geq \delta n^3 - 1$.

We have

$$\begin{aligned} \mathbb{E} \left(\sum_{m=1}^{b_n(l_n)} \pi_m^n g_m \right) &= \frac{1}{n^4} \sum_{l=0}^{l_n} \mathbb{E} \left(\sum_{m=a_n(l)}^{b_n(l)} g_m \right) \\ &\geq \frac{1}{n^4} \sum_{l \in \overline{M_n(l_n)}} \mathbb{E} \left(1_{\Omega_n(l)} \sum_{m=a_n(l)}^{b_n(l)} g_m \right). \end{aligned} \quad (15)$$

If $l \in \overline{M_n(l_n)}$ and $k_{a_n(l)} = \omega_1$, Player 2 plays L as long as $k_m = \omega_1$ and $m \leq b_n(l)$, while Player 1 plays $(1 - n^{-2}) \cdot T + n^{-2} \cdot B$. As a result, if $k_{a_n(l)} = \omega_1$, the probability that the state remains in ω_1 until stage $b_n(l)$ is $\alpha_n := (1 - n^{-2})^n$. Thus, the last inequality yields

$$\mathbb{E} \left(\sum_{m=1}^{b_n(l_n)} \pi_m^n g_m \right) \geq n^{-3} \text{Card } \overline{M_n(l_n)} \mathbb{P}(\Omega_n) \alpha_n \quad (16)$$

$$\geq (n^{-3}(l_n + 1) - \delta) \mathbb{P}(\Omega_n) \alpha_n. \quad (17)$$

Case 1 $l_n = n^3 - 1$.

By (16) and Lemma 1, there exists $\bar{n} \in \mathbb{N}^*$ such that for all $n \geq \bar{n}$ verifying $l_n = n^3 - 1$,

$$\gamma_{\pi^n}^{\omega_1}(\sigma^n, \tau^n) \geq 1 - 2\delta. \quad (18)$$

Case 2 $l_n < n^3 - 1$.

Let $n \in \mathbb{N}^*$ such that $l_n < n^3 - 1$. In particular, $|M_n(l_n)| \geq \delta n^3 - 1$ and $|\overline{M_n(l_n)}| \leq l_n - \delta n^3 + 2$.

We are going to show the following inequality:

$$\mathbb{P}(k_{b_n(l_n)} = 1^*) \geq \mathbb{P}(\Omega_n) - \left(1 - n^{-2} (1 - n^{-2})^n\right)^{\delta n^3 - 1} := \beta_n. \quad (19)$$

The idea is the following. Each time Player 2 plays R in state ω_1 , the state goes to 1^* with probability n^{-2} . If $k_{a_n(l)} = \omega_1$ and $l \in M_n(l)$, then at each stage $m \in \{a_n(l), a_n(l) + 1, \dots, m_n(l) - 1\}$, Player 2 will play L , hence at each of these stages the state will remain in ω_1 with probability n^{-2} . Since $m_n(l) - a_n(l) \leq n$,

with high probability $k_{m_n(l)} = \omega_1$. At stage $m_n(l)$, Player 2 plays R . Thus with high probability, conditionnal to the event $\Omega_n(l_n)$, before stage $b_n(l_n)$ Player 2 has played more than $\delta n^3 - 1$ times the action R in state ω_1 , leading the state to be absorbed in 1^* before stage $b_n(l_n)$ with high probability.

Formally, if $l \in \{0, 1, \dots, l_n\}$, we have

$$\begin{aligned} \mathbb{P}(\{k_{b_n(l)} \neq 1^*\} \cap \Omega_n(l)) &= \mathbb{P}(\{k_{b_n(l)} \neq 1^*\} \cap \{k_{a_n(l)} = \omega_1\} \cap \Omega_n(l)) \\ &= \mathbb{P}(\{k_{b_n(l)} \neq 1^*\} \mid \{k_{a_n(l)} = \omega_1\} \cap \Omega_n(l)) \\ &\quad \times \mathbb{P}(\{k_{a_n(l)} = \omega_1\} \cap \Omega_n(l)). \end{aligned} \quad (20)$$

First we majorize the first term $P_1 := \mathbb{P}(k_{b_n(l)} \neq 1^* \mid \{k_{a_n(l)} = \omega_1\} \cap \Omega_n(l))$. If $l \notin M_n(l)$, we simply majorize it by 1. Assume now that $l \in M_n(l)$. By the Markov property (σ^n and τ^n are Markovian strategies), we have

$$\begin{aligned} P_1 &= \mathbb{P}(k_{b_n(l)} \neq 1^* \mid \{k_{a_n(l)} = \omega_1\}) \\ &= \mathbb{P}(\{k_{b_n(l)} \neq 1^*\} \cap \{k_{m_n(l)} = \omega_1\} \mid \{k_{a_n(l)} = \omega_1\}) \\ &\quad + \mathbb{P}(\{k_{b_n(l)} \neq 1^*\} \cap \{k_{m_n(l)} \neq \omega_1\} \mid \{k_{a_n(l)} = \omega_1\}). \end{aligned}$$

Let $P_3 := \mathbb{P}(k_{m_n(l)} \neq \omega_1 \mid k_{a_n(l)} = \omega_1)$. The last equality and the Markov property give

$$\begin{aligned} P_1 &\leq \mathbb{P}(k_{b_n(l)} \neq 1^* \mid \{k_{m_n(l)} = \omega_1\} \cap \{k_{a_n(l)} = \omega_1\})(1 - P_3) + P_3 \\ &= \mathbb{P}(k_{b_n(l)} \neq 1^* \mid k_{m_n(l)} = \omega_1)(1 - P_3) + P_3. \end{aligned} \quad (21)$$

If $k_{m_n(l)} = \omega_1$, then at stage $m_n(l)$ Player 2 plays the action R , hence the state is absorbed in 1^* with probability n^{-2} . Thus

$$\mathbb{P}(k_{b_n(l)} \neq 1^* \mid k_{m_n(l)} = \omega_1) \leq 1 - n^{-2}. \quad (22)$$

If $k_{a_n(l)} = \omega_1$, then at each stage $m \in \{a_n(l), a_n(l) + 1, \dots, m_n(l) - 1\}$, Player 2 will play L , hence at each stage the state will remain in ω_1 with probability $1 - n^{-2}$, and $m_n(l) - a_n(l) \leq n$. We deduce that

$$P_3 \leq 1 - (1 - n^{-2})^n. \quad (23)$$

Combining (21), (22) and (23) gives

$$\begin{aligned} P_1 &\leq (1 - n^{-2})(1 - P_3) + P_3 \\ &= 1 + n^{-2}(P_3 - 1) \\ &\leq 1 - n^{-2}(1 - n^{-2})^n. \end{aligned} \quad (24)$$

As for the second term in (20), we have

$$\mathbb{P}(\{k_{a_n(l)} = \omega_1\} \cap \Omega_n(l)) \leq \mathbb{P}(\{k_{b_n(l-1)} \neq 1^*\} \cap \Omega_n(l-1)). \quad (25)$$

Combining (20), (24) and (25), we deduce that if $l \in M_n(l)$, then

$$\mathbb{P}(\{k_{b_n(l)} \neq 1^*\} \cap \Omega_n(l)) \leq (1 - n^{-2}(1 - n^{-2})^n) \mathbb{P}(\{k_{b_n(l-1)} \neq 1^*\} \cap \Omega_n(l-1)).$$

Because $|M_n(l_n)| \geq \delta n^3 - 1$, by induction we obtain

$$\mathbb{P}(\{k_{b_n(l_n)} \neq 1^*\} \cap \Omega_n(l_n)) \leq \left(1 - n^{-2} (1 - n^{-2})^n\right)^{\delta n^3 - 1},$$

and inequality (19) follows. Now we can minorize the other part of the payoff:

$$\begin{aligned} \mathbb{E} \left(\sum_{m \geq b_n(l_n)+1} \pi_m^n g_m \right) &\geq \mathbb{E} \left(1_{\{k_{b_n(l_n)} = 1^*\}} \sum_{m \geq b_n(l_n)+1} \pi_m^n \right) \\ &= n^{-3} (n^3 - l_n - 1) \mathbb{P}(\{k_{b_n(l_n)} = 1^*\}) \\ &\geq (1 - n^{-3} (l_n + 1)) \beta_n. \end{aligned} \quad (26)$$

Inequalities (17) and (26) yield

$$\begin{aligned} \mathbb{E} \left(\sum_{m \geq 1} \pi_m^n g_m \right) &= \mathbb{E} \left(\sum_{m=1}^{b_n(l_n)} \pi_m^n g_m \right) + \mathbb{E} \left(\sum_{m \geq b_n(l_n)+1} \pi_m^n g_m \right) \\ &\geq (n^{-3} (l_n + 1) - \delta) \mathbb{P}(\Omega_n) \alpha_n + (1 - n^{-3} (l_n + 1)) \beta_n. \end{aligned}$$

The sequences $(\alpha_n)_{n \geq 1}$, $(\beta_n)_{n \geq 1}$ and $(\mathbb{P}(\Omega_n))_{n \geq 1}$ converge to 1, thus there exists $n_1 \in \mathbb{N}^*$ such that for all $n \geq n_1$ verifying $l_n < n^3 - 1$, we have

$$v_{\pi^n}(\omega_1) \geq \gamma_{\pi^n}^{\omega_1}(\sigma^n, \tau^n) \geq 1 - 2\delta. \quad (27)$$

Because τ^n is a best-response strategy to σ^n in $\Gamma_{\pi^n}^{\omega_1}$, inequalities (18) and (27) show that for $n \geq \max(\bar{n}, n_1)$, we have

$$v_{\pi^n}(\omega_1) \geq \gamma_{\pi^n}^{\omega_1}(\sigma^n, \tau^n) \geq 1 - 2\delta. \quad (28)$$

Because $\delta \in (0, 1]$ is arbitrary, the sequence $(v_{\pi^n}(\omega_1))_{n \geq 1}$ converges to 1, and Γ has no 1-asymptotic value. \square

Now we can prove Theorem 4.

Proof (Proof of Theorem 4) Let $\epsilon > 0$ and $p := 1/2 + \epsilon$. Proving that Γ has no p -asymptotic value proceeds in the same way as previously. The only difference is that the sequence of weights (π^n) has to be modified. Let $\epsilon > 0$ and $n \in \mathbb{N}^*$. In what follows, the integer part of a real number x is denoted by $\lfloor x \rfloor$. Define two integers N_1 and N_2 by

$$N_1 := \lfloor n^{2-\epsilon} \rfloor \quad \text{and} \quad N_2 := \lfloor n^{2+\epsilon} \rfloor.$$

For $l \in \{0, 1, \dots, N_2 - 1\}$, let $a'_n(l) := l(N_1 + n^5) + 1$ and $b'_n(l) := l(N_1 + n^5) + N_1$. Let

$$E'_1 := \bigcup_{l \in \{0, 1, \dots, N_2\}} \{a'_n(l), a'_n(l) + 1, \dots, b'_n(l)\}.$$

Let $\pi'^n \in \Delta(\mathbb{N}^*)$ defined in the following way: for $m \in \mathbb{N}^*$,

$$\pi'_m := \begin{cases} n^{-4} & \text{if } m \in E'_1 \setminus \{N_1 + 1\}, \\ 1 - \sum_{m \neq N_1+1} \pi'_m & \text{if } m = N_1 + 1, \\ 0 & \text{if } m \notin E'_1. \end{cases}$$

We have

$$I_p(\pi'^n) \leq \lfloor n^{2+\epsilon} \rfloor n^{-4(1/2+\epsilon)} + 2\pi'_{N_1+1}.$$

Hence $\lim_{n \rightarrow +\infty} I_p(\pi'^n) = 0$. We claim that $\lim_{n \rightarrow +\infty} v_{\pi'^n}(\omega_1) = 1$. The proof is the same as above. We still consider the same strategy σ^n for Player 1 in $\Gamma^{\pi'^n}$. Lemma 1 is still true. Indeed, the length of the blocks $\{b'_n(l) + 1, b'_n(l) + 2, \dots, a'_n(l + 1) - 1\}$ is still n^5 .

Now let us check the remainder of the proof. The quantity $(1 - n^{-2})^{n^{2+\epsilon}}$ goes to 1 as n goes to infinity. Hence if $l \in \{0, \dots, N_2\}$ and $k_{l(N_1+n^5)+1} = \omega_1$, to get a good payoff between stage $a'_n(l)$ and stage $b'_n(l)$, Player 2 should make the state move from ω_1 to ω_2 at least before stage $b_n(l)$. Thus, he has to play R at least one time, and take a risk of being absorbed in ω_1 of n^{-2} . There are approximately $n^{2+\epsilon}$ such blocks. Since $(1 - n^{-2})^{n^{2+\epsilon}}$ goes to 0 as n goes to infinity, the same proof as before shows that, when n goes to infinity, the sequence $(v_{\pi^n}(\omega_1))_{n \geq 1}$ converges to 1. \square

3 Uniform approach

To relax the assumption that sequences of weights are decreasing in Corollary 1, the simplest sequences of weights one can imagine are the $\pi^{l,n}$ defined in Remark 3: $\pi^{l,n} := n^{-1} 1_{l+1 \leq m \leq l+n}$. As we have seen, Theorem 3 shows that for any stochastic game,

$$\lim_{n \rightarrow +\infty} \sup_{l \in \mathbb{N}} \|v_{\pi^{l,n}} - v_\infty\|_\infty = 0.$$

Is it possible to show the existence of strategies that are approximately optimal in any game $\Gamma_{\pi^{l,n}}$, for any $l \geq 0$ and n big enough, for both players? We provide an example of an absorbing game where this property does not hold. Thus, no natural extension of Theorem 1 to sequences of weights which are not decreasing seems to exist.

Consider the following absorbing game, introduced by Gillette [1957] under the name of "Big Match". The state space is $K = \{\omega, 1^*, 0^*\}$, where 1^* (resp. 0^*) is an absorbing state with payoff 1 (resp. 0). Action sets for Player 1 and 2 are respectively $I = \{T, B\}$ and $J = \{L, R\}$. The payoff and transition functions in state ω are described by the following table:

Fig. 1 Transition and payoff functions in state ω

J1 \ J2	L	R
T	1*	0*
B	0	1

As any stochastic game, the Big Match has a uniform value v_∞ , and $v_\infty(\omega) = 1/2$ (see [Sorin, 2002, Chapitre 5, p.93]). The stationary strategy $1/2 \cdot L + 1/2 \cdot R$ is a 0-optimal uniform strategy for Player 2. Constructing ϵ -optimal uniform strategy for Player 1 is more tricky (see Blackwell and Ferguson [1968]).

Now we investigate the general uniform approach in the Big Match.

Definition 6 Let Γ be a stochastic game, and k_1 the initial state. Player 1 can guarantee uniformly *in the general sense* $\alpha \in \mathbb{R}$ in Γ^{k_1} if for all $\epsilon > 0$, there exists $\sigma^* \in \Sigma$ and $N_1 \in \mathbb{N}^*$ such that for all $\tau \in \mathcal{T}$, $n \geq N_1$ and $\bar{n} \in \mathbb{N}$, we have

$$\mathbb{E}_{\sigma^*, \tau}^{k_1} \left(\frac{1}{n} \sum_{m=\bar{n}+1}^{\bar{n}+n} g_m \right) \geq \alpha - \epsilon. \quad (29)$$

First, let us explain why Player 1 can not guarantee uniformly more than 0 in the general sense. Assume the contrary: Player 1 can guarantee uniformly in the general sense $\alpha > 0$. Let $(N_1, \sigma^*) \in \mathbb{N}^* \times \Sigma$ corresponding to $\epsilon = \alpha/2$ in (29). The stationary strategy $\alpha/10 \cdot L + (1 - \alpha/10) \cdot R$ is denoted by y .

On the one hand, Player 1 should not play T at any stage of the game, against the strategy y . On the other hand, in the infinitely repeated game, with high probability Player 2 will play L at N_1 random consecutive stages. At that point, Player 1 does not know if Player 2 has switched to a pure strategy that plays L at any stage, or if he still plays y . In the first case, Player 1 should play T at least one time during these N_1 stages, but in the second case, he should not play T . Thus, he cannot guarantee a good payoff against both strategies. This provides intuition for the following proposition:

Proposition 4 *Player 1 cannot guarantee uniformly more than 0 in the general sense.*

Remark 6 This proposition shows in particular that Theorem 2 (Renault and Venel) does not generalize to zero-sum stochastic games, even for absorbing games: the game Γ has no general uniform value.

Proof The same notations as in the above discussion are used.

For $n \in \mathbb{N}^*$, let A_n be the event {Player 1 plays T before stage n }, and let \bar{A}_n be the complement of A_n . The sequence $(\mathbb{P}_{\sigma^*, y}(A_n))_{n \geq 1}$ is increasing and bounded by 1, therefore it converges to some $l \in [0, 1]$. Let $N_0 \in \mathbb{N}^*$ such that for all $n \geq N_0$,

$$\mathbb{P}_{\sigma^*, y}(A_n) \geq l - \alpha/10.$$

To avoid confusion, in what follows h denotes an element of H_∞ , and \tilde{h} denotes the random variable with values in H_∞ describing the infinite history of the game. Let $n \geq N_0 + N_1$. Let $H^n \subset H_\infty$ be the following set:

$$\{h \in H_\infty \mid \exists a(h) \in \{N_0, \dots, n - N_1\}, \forall m \in \{a(h) + 1, \dots, a(h) + N_1\}, j_m = L\}.$$

There exists $N_2 \in \mathbb{N}^*$ such that

$$\mathbb{P}_{\sigma^*, y}(\tilde{h} \in H^{N_2}) \geq 1/2. \quad (30)$$

We have

$$\mathbb{P}_{\sigma^*, y}(\overline{A_{N_0}} \cap A_{N_2}) \geq \mathbb{E}_{\sigma^*, y} \left(1_{\{\tilde{h} \in H^n\}} \mathbb{P}_{\sigma^*, y'(\tilde{h})}(\overline{A_{N_0}} \cap A_{N_2}) \right), \quad (31)$$

where for $h \in H^n$, the strategy $y'(h)$ is the Markov strategy equal to y between stages 1 and $a(h)$, and equal to $j_m(h)$ for each stage $m \geq a(h) + 1$.

Let $h \in H^n$. Let us now minorize $\mathbb{P}_{\sigma^*, y'(h)}(\overline{A_{N_0}} \cap A_{N_2})$.

If Player 1 plays T at some stage $m \leq a(h)$ against the strategy $y'(h)$, the game is absorbed in 1^* with probability $\alpha/10$, and in 0^* with probability $1 - \alpha/10$. Therefore we have

$$\mathbb{E}_{\sigma^*, y'(h)} \left(1_{A_{N_0}} \frac{1}{N_1} \sum_{m=a(h)+1}^{a(h)+N_1} g_m \right) \leq \frac{\alpha}{10}. \quad (32)$$

Because $h \in H^n$, we have

$$\mathbb{E}_{\sigma^*, y'(h)} \left(1_{\overline{A_{N_2}}} \frac{1}{N_1} \sum_{m=a(h)+1}^{a(h)+N_1} g_m \right) = 0. \quad (33)$$

Combining (32) and (33), we obtain

$$\mathbb{E}_{\sigma^*, y'(h)} \left(\frac{1}{N_1} \sum_{m=a(h)+1}^{a(h)+N_1} g_m \right) \leq \frac{\alpha}{10} + \mathbb{P}_{\sigma^*, y'(h)}(\overline{A_{N_0}} \cap A_{N_2}),$$

and by (29),

$$\mathbb{P}_{\sigma^*, y'(h)}(\overline{A_{N_0}} \cap A_{N_2}) \geq \alpha/2 - \alpha/10 = 2\alpha/5.$$

Plugging the last inequality into (31) and using (30), we deduce that

$$\begin{aligned} \mathbb{P}_{\sigma^*, y}(\overline{A_{N_0}} \cap A_{N_2}) &\geq \mathbb{E}_{\sigma^*, y} \left(1_{\{h \in H^n\}} 2\alpha/5 \right) \\ &\geq \alpha/5. \end{aligned}$$

Because $A_{N_0} \subset A_{N_2}$, we have

$$\mathbb{P}_{\sigma^*, y}(\overline{A_{N_0}} \cap A_{N_2}) = \mathbb{P}_{\sigma^*, y}(A_{N_2}) - \mathbb{P}_{\sigma^*, y}(A_{N_0}) \leq \alpha/10,$$

thus $\alpha/10 \geq \alpha/5$, which is a contradiction. \square

Acknowledgments

I would like to thank Jérôme Renault, Sylvain Sorin, Rida Laraki, and Fabien Gensbittel for their interesting suggestions.

References

- T. Bewley and E. Kohlberg. The asymptotic theory of stochastic games. *Mathematics of Operations Research*, 1(3):197–208, 1976.
- D. Blackwell and T. Ferguson. The big match. *The Annals of Mathematical Statistics*, 39(1):159–163, 1968.
- P. Cardaliaguet, R. Laraki, and S. Sorin. A continuous time approach for the asymptotic value in two-person zero-sum repeated games. *SIAM Journal on Control and Optimization*, 50(3):1573–1596, 2012.
- Gillette. Stochastic games with zero stop probabilities. In *Contributions to the Theory of Games*, pages 179–187. Princeton University Press, 1957.
- A. Maitra and T. Parthasarathy. On stochastic games. *Journal of Optimization Theory and Applications*, 5(4):289–300, 1970.
- J. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10(2):53–66, 1981.
- J.-F. Mertens, A. Neyman, and D. Rosenberg. Absorbing games with compact action spaces. *Mathematics of Operations Research*, 34(2):257–262, 2009.
- A. Neyman and S. Sorin. Repeated games with public uncertain duration process. *International Journal of Game Theory*, 39(1):29–52, 2010.
- M. Oliu-Barton. The asymptotic value in finite stochastic games. *Mathematics of Operations Research*, 39(3):712–721, 2014.
- J. Renault and X. Venel. A distance for probability spaces, and long-term values in markov decision processes and repeated games. *Arxiv preprint arXiv:1202.6259*, 2012.
- L. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39(10):1095, 1953.
- S. Sorin. *A first course on zero-sum repeated games*, volume 37. Springer, 2002.
- G. Vigeral. A zero-sum stochastic game with compact action sets and no asymptotic value. *Dynamic Games and Applications*, 3(2):172–186, 2013.
- B. Ziliotto. Zero-sum repeated games: counterexamples to the existence of the asymptotic value and the conjecture $\max\min = \lim v(n)$. *arXiv preprint arXiv:1305.4778*, to appear in *Annals of Probability*, 2013.